AD-A044 897 STANFORD UNIV CALIF DEPT OF OPERATIONS RESEARCH F/G 12/2

OPTIMAL SELLING WHEN THE PRICE DISTRIBUTION IS UNKNOWN.(U)

SEP 77 C DERMAN, G J LIEBERMAN, S ROSS N00014-75-C-0561

UNCLASSIFIED TR-185 NL

| OF |

AD

A044 897

END

DATE

FILMED

10-77

DDC

6 OPTIMAL SELLING WHEN THE PRICE DISTRIBUTION IS UNKNOWN,

by

10 C. DERMAN, G. J. LIEBERMAN and S. ROSS

9 TECHNICAL REPORT NO. 185

11 6 SEPTEMBER 6, 1977 12 26p.

14 TR-185

15 SUPPORTED BY THE ARMY AND NAVY
~~UNDER CONTRACT~~ N00014-75-C-0561 (NR 042-002)
WITH THE OFFICE OF NAVAL RESEARCH

Gerald J. Lieberman, Project Director

DEPARTMENT OF OPERATIONS RESEARCH
AND
DEPARTMENT OF STATISTICS
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

1473
402 766

LB

OPTIMAL SELLING WHEN THE PRICE DISTRIBUTION IS UNKNOWN[*]

by

C. Derman, Columbia University
C. J. Lieberman, Stanford University
S. M. Ross, University of California, Berkeley

0. Introduction

This paper reconsiders the classical model for selling an asset in which offers come in daily and a decision must then be made as to whether or not to sell. For each day the item remains unsold a continuation (or maintenance cost) c is incurred. The successive offers are assumed to be independent and identically distributed random variables having an unknown distribution F. The model is considered both in the case where c[illegible] an offer is rejected it may not be recalled at a later ti[illegible] in the case where such recall of previous offers is allowed.

In Section 1 we show how bounds on the optimal policy may be obtained when some partial information about F is available. In particular, we show that if F, the distribution of offers, satisfies the NWUE (new worse than used in expection) property defined as

$$E_F[X - a \mid X > a] \geq E_F[X] \qquad \text{for all } a \geq 0,$$

[*] This research was supported in part by the Office of Naval Research under contracts N00014-75-C-0620 with Columbia University, N00014-75-C-0561 with Stanford University, and N00014-68-A-0200-1036 with the University of California, and by the U.S. Army Research Office-Durham under contract DAHC04-74-C-0226 with the University of California.

then the optimal policy has a monotonic relationship with the optimal policy in the case where the distribution of offers is exponential with the same mean as F.

In Sections 2 and 3 we consider a Bayesian version of this model by supposing that F is known to be one of the distributions $F_1, F_2, \ldots, F_n$ with given initial prior probabilities. In Section 2 we do not allow, and in Section 3 we do allow, the recall of old offers. In both cases we provide bounds on the optimal policy in terms of the optimal policies in the case where it is known which of the F_i is equal to F. This Bayesian format has previously been considered in [2] which assumed that F was a normal random variable with known variance and imposed a normal prior distribution on the mean of F. As our model imposes no parametric condition on F in the prior distribution, the type of results we obtain are somewhat different than those in [2].

1. Independent and Identically Distributed Offers from an Unknown Distribution with Partial Information

If the successive offers were independent and identically distributed random variables having known distribution F, then it is

well-known (see [1]) that the policy that maximizes the total expected return, both with and without recall, is to accept an offer x if and only if $x \geq x_F$, where x_F is the smallest value such that

$$x_F \geq \left[\int_{x_F}^{\infty} x dF(x) - c\right] \Big/ (1 - F(x_F)) \quad .$$

If F is continuous, this reduces to

$$c = \int_{x_F}^{\infty} (x - x_F)\, dF(x) \; .$$

The optimal expected return is $x_F + c$.

We shall start out by comparing the optimal critical number for two different distributions. To begin we need the following definition.

Definition:

For any two probability distributions F and G we say that

$F \underset{V}{<} G$ if

$$\int f(x)\, dF(x) \leq \int f(x)\, dG(x)$$

for all increasing convex functions f.

If F and G have the same means, then $F \underset{V}{<} G$ intuitively means that F has less variability than G.

Proposition 1:

Let $F \underset{V}{\leq} G$ then $x_F \leq x_G$.

Proof:

x_G is the smallest value satisfying

$$c \geq E_G[(X - x_G)^+] \ .$$

Now

$$E_G[(X - x_G)^+] \geq E_F[(X - x_G)^+]$$

since $f(x) = (x - x_G)^+$ is an increasing convex function.

Hence,

$$c \geq E_F[(X - x_G)^+]$$

implying that $x_F \leq x_G$. ||

Proposition 2 is concerned with the return from a nonoptimal policy.

Proposition 2:

If $x \leq x_F$, then the return from the policy that accepts the first offer that is at least as large as x has a return that is at least $x + c$.

Proof:

To prove the above, consider the expected difference between the optimal policy that uses the critical number x_F and the above policy that uses the critical number x. By conditioning on whether an offer between x and x_F occurs before or after an offer greater than x_F, we see that the expected difference is at most $x_F - x$ in the former case (since the expected return from the optimal policy starting at the time of this offer between x and x_F is equal to $x_F + c - c = x_F$) and it is 0 in the latter case. Hence, the result follows. ||

Definition:

We say that the distribution F, with $F(0-) = 0$, is NWUE if

$$\int_a^\infty \frac{\bar{F}(x)\,dx}{\bar{F}(a)} \geq \int_0^\infty \bar{F}(x)\,dx \qquad \text{for all } a \geq 0 \text{ ,}$$

where $\bar{F}(x) = 1 - F(x)$. (If X is a random variable having distribution F, then the above is equivalent to $E[X - a | X > a] \geq E[X]$.)

Proposition 3:

If F is NWUE with mean μ, then

$$E(\mu) \underset{V}{\leq} F$$

where $E(\mu)$ is an exponential distribution with mean μ.

Proof:

It is easy to show that $F \underset{V}{\geq} G$ is equivalent to

$$\int_a^\infty \bar{F}(x)\, dx \geq \int_a^\infty \bar{G}(x)\, dx \quad \text{for all } a .$$

Thus, we have to show that

$$\int_a^\infty \bar{F}(x)\, dx \geq \mu e^{-a/\mu}$$

whenever F is NWUE. By the definition of NWUE we have

$$\int_a^\infty \frac{\bar{F}(x)}{\mu}\, dx \geq \bar{F}(a) ,$$

or equivalently,

$$\bar{F}_e(a) \geq \mu \frac{\bar{F}(a)}{\mu} = \mu f_e(a) ,$$

where

$$\bar{F}_e(a) = \int_a^\infty \frac{\bar{F}(x)}{\mu}\, dx \quad \text{and} \quad f_e(a) = -\frac{d\, F_e(a)}{da}$$

are the equilibrium distribution corresponding to F. Hence,

$$\frac{f_e(a)}{\bar{F}_e(a)} \le \frac{1}{\mu} \ ,$$

thereby implying, upon integrating, that

$$-\log \bar{F}_e(a) \le a/\mu \ ,$$

which proves the result. ||

We are now ready for the main theorem of this section.

Theorem 1:

If the unknown distribution F is known to be NWUE and to have mean μ then

$$x_F \ge \bar{x}$$

and the policy which accepts the first offer of at least $\bar{x}$ has return of at least $\bar{x} + c$; where

$$\bar{x} = -\mu \log(c/\mu) \ .$$

Proof:

The result follows immediately from Propositions 1, 2, and 3 since $\bar{x} = x_E$ when E is an exponential distribution with mean μ. ||

Remark:

One instance in which the distribution of offers would be NWUE is if there were many classes of potential customers and offers from each class followed an exponential distribution. Thus the distribution of offers would be a mixture of exponential distributions and the degenerate distribution at 0 (indicating no offer), and would thus be NWUE.

2. A Bayesian Model Without Recall of Past Offers

In this section we suppose that if an offer is rejected then it can never be accepted in the future. In addition, we suppose that although the distribution F is not known with certainty, we do know that is is one of the distributions $F_1, F_2, \ldots, F_n$, with given prior probabilities. We say that the state of the system is $(x, \underline{P})$ when x is the present offer under consideration and $\underline{P} = (P_1, \ldots, P_n)$ is the posterior probability vector, given all the information that we have accumulated up to that point (including the present offer x), as to which of the F_i is the actual distribution.

Also define $V(x, \underline{P})$ to equal the expected return from this day onward given that the state today is $(x, \underline{P})$ and we employ an optimal policy. (If we assume, as we do, that each of the F_i has a finite variance and $c > 0$ then it can be shown as in [1] that an optimal policy exists).

The optimality equation thus takes the following form

$$V(x, \underline{P}) = \mathrm{Max}\{x, V(\underline{P}) - c\} ,$$

where $V(\underline{P})$, which represents the best you can do when the distribution is chosen by the prior probability vector $\underline{P} = (P_1, \ldots, P_n)$, satisfies

$$V(\underline{P}) = \Sigma P_j \int V(y, T_y \underline{P}) \, dF_j(y)$$

where

$$T_y \underline{P} = ((T_y \underline{P})_1, \ldots, (T_y \underline{P})_n)$$

and

$$(T_y P)_j = \text{Prob}\{F_j | \underline{P}, y\}$$

$$= \frac{P_j \, dF_j(y)}{\Sigma P_i \, dF_i(y)} .$$

Furthermore, the optimal policy accepts the offer in state $(x, \underline{P})$ if and only if

$$x \geq V(\underline{P}) - c .$$

Proposition 4:

$V(\underline{P})$ is convex function of $\underline{P}$.

Proof:

Recall that $V(\underline{P})$ represents the best we can do when the distribution is chosen according to $\underline{P}$. Now suppose $\underline{P} = \lambda \underline{P}^1 + (1 - \lambda) \underline{P}^2$, for some $0 < \lambda < 1$, and suppose that the distribution to be used is to be chosen according to the following two-stage experiment. First we flip a coin having probability λ of coming up heads. If the coin comes up heads then we choose the distribution according to the prior probability $\underline{P}^1$, while if it comes up tails then we use $\underline{P}^2$. Now if

we are not told the outcome of the coin flip then the problem is exactly the same as if the distribution was chosen according to $\underline{P}$ and thus, the best we can do is $V(\underline{P})$. On the other hand, if we are to be told about the outcome of the coin flip then by conditioning on the outcome we see that our expected return if we play optimally is $\lambda V(\underline{P}^1) + (1 - \lambda) V(\underline{P}^2)$. Hence, as additional information can not lower our expected return we see that

$$V(\underline{P}) \leq \lambda V(\underline{P}^1) + (1 - \lambda) V(\underline{P}^2)$$

and the result is proven. ||

Let $x_i = x_{F_i}$, $i = 1, \ldots, n$.

Corollary 1:

$V(\underline{P}) \leq \Sigma P_i x_i + c.$

Proof:

This follows directly from Proposition 4 since $V(0,0,0,1,0\ldots0) = x_i + c$ (where the 1 is in the ith place).

Proposition 5:

If the present state is $(x, \underline{P})$ then

(i) if $x \geq \Sigma P_i x_i$ then it is optimal to accept x,

(ii) if

$$x < \sum_i P_i \left[\frac{\int_x^{\infty} y \, dF_i(y) - c}{1 - F_i(x)} \right]$$

then it is optimal to reject the offer x,

(iii) if

$$x < \sum_i P_i \int_{-\infty}^{\infty} y dF_i(y) - c$$

then it is optimal to reject x.

Proof:

(i) If $x \geq \Sigma P_i x_i$ then, using Corollary 1, we have

$$x \geq V(P) - c$$

and (i) is established.

(ii) Suppose the present state is $(x, \underline{P})$ and consider the policy that accepts the first offer greater than x. The expected return from this policy is

$$\sum_i P_i \left[\int_x^{\infty} \frac{ydF_i(y)}{1-F_i(x)} - \frac{c}{1-F_i(x)} \right]$$

which follows by noting that, given that the distribution is F_i, then the expected number of additional offers that will be

made until one is accepted is $1/(1-F_i(x))$. Clearly if x is less than this value, then it cannot be optimal to accept the present offer of x.

(iii) The proof of (iii) is similar to that of (ii) in that it considers the return when in state $(x, \underline{P})$ if you accept the next offer, and notes that if this return is greater than x then x should clearly not be accepted. ||

Remark:

It follows from part (ii) of the above proposition that if $x < \min(x_1, \ldots, x_n)$ then it is always optimal to reject x. ||

Let us now consider the special case where there are only two possible distributions, i.e., F_1 and F_2 and suppose $x_1 \leq x_2$. In this case the state can be represented as the pair (x, P) where x is the present offer and P is the present probability (given all information, including x, accumulated up to this point) that F_2 is the true distribution. In this case we have

Theorem 2:

$V(P)$ is an increasing function of P, $0 \leq P \leq 1$.

Proof:

Since $V(P)$ is a convex function of P(Proposition 4) the result would follow if we could show that

$$V(0) \leq V(P) \quad \text{for all} \quad 0 \leq P \leq 1 .$$

Now $V(0) = x_{F_1} + c \equiv x_1 + c$. Also, as it is always optimal to reject an offer less than $\min(x_1, x_2) = x_1$ it follows from the optimality equation that

$$V(P) - c \geq x_1 \qquad \text{for all} \quad P$$

which proves the result. ||

Thus, when $n = 2$ and $x_1 \leq x_2$, it is optimal to accept the offer when in state (x, P) if and only if $x \geq h(P)$ where $h(P) \equiv V(P) - c$ is an increasing convex function of P with $h(0) = x_1$, $h(1) = x_2$. Furthermore, bounds on $h(P)$ are given by Proposition 5.

Remark:

There does not appear to be an analogue to Theorem 2 when there are more than 2 possible distributions. For instance, suppose that the distributions $F_1, F_2, \ldots, F_n$ are stochastically increasing in the sense that $F_i(t)$ is nonincreasing in i for each t. If we define the probability vector $\underline{P}$ to be greater than or equal to the probability vector $\underline{Q}$, written $\underline{P} \geq \underline{Q}$, if

$$\sum_{1}^{j} P_i \leq \sum_{1}^{j} Q_i \qquad \text{for each} \quad j = 1, \ldots n ,$$

then we might hope to prove that $V(\underline{P}) \geq V(\underline{Q})$. However, this need not be

the case as is indicated by the following example. Suppose F_1 puts all its weight on the value .9, F_2 puts all its weight on the value 1, and F_3 is the distribution of a random variable that takes on the value 1 with probability .99 and $(10)^6$ with probability .01, and suppose $c = 1$. Now, $\underline{P} \equiv (0, .9, .1) \geq \underline{Q} \equiv (.9, 0, .1)$, but it turns out that $V_{(\underline{P})} < V_{(\underline{Q})}$; the reason being that under $\underline{Q}$ it only takes a single observation to determine the true F_i, whereas this is not so under $\underline{P}$. ||

3. Independent and Identically Distributed Offers from an Unknown Distribution with Recall of Past Offers

In the previous section we assumed that once an offer was rejected by the decision maker, then that offer immediately disappears. In this section, however, we consider the same model as in Section 2 but with the exception that an offer remains good indefinitely and may be accepted at any time.

It turns out that when the distribution of offers is known, then the optimal policy in this case is identical to the one where recalling past offers is not allowed. That is, the optimal policy is to accept the first offer that is at least as large as x_F and the expected return under the optimal policy is $x_F + c$, when x_F is as defined in Section 1.

Consider now the case where the distribution of offers is one of the distributions $F_1, \ldots, F_n$, where the F_i is chosen according to some initial probability vector. The state of the system at any time can be defined by $(m, \underline{P})$ where m is the maximum offer that has been received up to that time and $\underline{P}$ is the posterior probability vector (given all offers up to that time, including any just made) of the true distribution. The optimality equation takes the form

$$V(m, \underline{P}) = \text{Max}\{m, \sum_i P_i \left[\int_0^m V(m, T_y\underline{P})\, dF_i(y) + \int_m^\infty V(y, T_y\underline{P})\, dF_i(y)\right] - c\} ,$$

where

$$T_y \underline{P} = ((T_y \underline{P})_1, \ldots, (T_y \underline{P})_n) ,$$

and

$$(T_y \underline{P})_j = \frac{P_j dF_j(y)}{\Sigma P_i dF_i(y)} .$$

While it follows from its definition that $V(m, \underline{P})$ is an increasing function of m for fixed $\underline{P}$ it is not immediately evident from the optimality equation that if the offer m is accepted when in state $(m, \underline{P})$ then the offer m^1 is also accepted when in state $(m^1, \underline{P})$ whenever $m^1 \geq m$. We now prove this.

Proposition 6:

For fixed $\underline{P}$, $V(m, \underline{P}) - m$ is a nonincreasing function of m.

Proof:

Suppose $m_1 < m_2$. Note that the distribution of the sequence of future offers is the same no matter whether the initial state is $(m_1, \underline{P})$ or $(m_2, \underline{P})$ (since it only depends on $(x, \underline{P})$ through $\underline{P}$). We can then conclude that if the initial state is $(m_1, \underline{P})$ then by following throughout the optimal policy for the initial state $(m_2, \underline{P})$ that our return when we stop is within $m_2 - m_1$ of what it would have been if the initial state was really $(m_2, \underline{P})$. Therefore,

$$V(m_1, \underline{P}) + m_2 - m_1 \geq V(m_2, \underline{P}) \; . \qquad ||$$

Corollary 2:

If it is optimal to accept m_1 when in state $(m_1, \underline{P})$ then it is optimal to accept m_2 when in state $(m_2, \underline{P})$ whenever $m_2 \geq m_1$.

Proof:

If $V(m_1, \underline{P}) = m_1$ then from Proposition 6

$$V(m_2, \underline{P}) - m_2 \leq 0 \; .$$

This implies, from the optimality equation, that

$$V(m_2, \underline{P}) = m_2 \; . \qquad ||$$

Proposition 7:

For fixed m, $V(m, \underline{P})$ is a convex function of $\underline{P}$.

Proof:

The proof is identical to the proof of Proposition 4 in in Section 2. ||

Corollary 3:

$$V(m, \underline{P}) \leq \Sigma P_i \operatorname{Max}(m, x_i)$$

where $x_i \equiv x_{F_i}$.

Proof:

Letting e_i be the vector of zero's with a one in the ith place then

$$V(m, e_i) = \begin{cases} m & \text{if } m > x_i \\ x_i & \text{if } m < x_i . \end{cases}$$

Hence, from convexity

$$V(m, \underline{P}) \leq \Sigma P_i V(m, e_i) = \sum_i P_i \operatorname{Max}\{m, x_i\} . \qquad ||$$

Proposition 8:

If the present state is $(m, \underline{P})$ then

(i) if $m > \Sigma P_i \operatorname{Max}(m, x_i)$ then it is optimal to accept m.

(ii) if

$$m < \sum_i P_i \left[\frac{\int_m^\infty y dF_i(y) - c}{1 - F_i(m)} \right]$$

then it is optimal to look at another offer

(iii) if

$$m < \sum_i P_i \left[mF_i(m) + \int_m^\infty y dF_i(y) \right] - c$$

then it is optimal to look at another offer.

Proof:

Part (i) follows directly from Corollary 3, while the proofs of parts (ii) and (iii) are identical to their corresponding results in Proposition 5 of Section 1. ||

Suppose now that $n = 2$ and $x_1 \leq x_2$. In this case we represent the state by (m, P) when P is the posterior probability that F_2 is the true distribution.

Theorem 3:

$V(m, P)$ is increasing in P for fixed m.

Proof.

As in the corresponding proof of the previous section we need to show that

$$V(m, 0) \leq V(m, P) .$$

Now,

$$V(m, 0) = \text{Max}(m, x_1) .$$

However,

$$V(m, P) \geq m ,$$

and as it follows from Part (ii) of Proposition 8 that it is never optimal to accept an offer less than x_1, we have

$$V(m, P) \geq x_1 .$$

That is,

$$m > x_1 \Rightarrow V(m, P) \geq x_1$$

$$m < x_1 \Rightarrow V(m, P) = V(x_1, P) > x_1 ,$$

and the proof is complete. ||

Hence, when $n = 2$ and $x_1 \leq x_2$, it is optimal to accept m when in state (m, P) if and only if $m \geq m(P)$, where $m(P)$ is an increasing convex function of P with $m(0) = x_1$, $m(1) = x_2$.

REFERENCES

[1] Chow, Y.S. and H. Robbins, "A Martingale Systems Theorem and Applications", Proc. 4th. Berkeley Symp. Math. Stat. and Prob., 1, 1961, 93-104.

[2] DeGroot, M.H., "Some Problems of Optimal Stopping", J. of the Royal Statistical Society", Series B, Vol. 30, No. 1, 1968, 108-122.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 185	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) OPTIMAL SELLING WHEN THE PRICE DISTRIBUTION IS UNKNOWN		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) C. Derman, G.J. Lieberman, and S. Ross		8. CONTRACT OR GRANT NUMBER(s) N00014-75-C-0561
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of Operations Research and Dept. of Statistics -- Stanford University Stanford, California 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS (NR-042-002)
11. CONTROLLING OFFICE NAME AND ADDRESS Statistics & Probability Program Code 436 Office of Naval Research Arlington, Virginia 22217		12. REPORT DATE September 6, 1977
		13. NUMBER OF PAGES 22
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

OPTIMAL SELLING, SELLING AN ASSET, OFFERS, BAYESIAN VERSION OF OPTIMAL SELLING, NWUE, RECALL OF PAST OFFERS, NON-RECALL OF PAST OFFERS.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This paper reconsiders the classical model for selling an asset in which offers come in daily and a decision must then be made as to whether or not to sell. For each day the item remains unsold a continuation(or maintenance cost) c is incurred. The successive offers are assumed to be independent and identically distributed random variables having an unknown distribution F.

over →

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

The model is considered both in the case where once an offer is rejected it may not be recalled at a later time, and in the case where such recall of previous offers is allowed.

In Section 1 we show how bounds on the optimal policy may be obtained when some partial information about F is available. In particular, we show that if F, the distribution of offers, satisfies the NWUE (new worse than used in expection) property defined as

$$E_F[X - a \mid X > a] \geq E_F[X] \qquad \text{for all } a \geq 0,$$

then the optimal policy has a monotonic relationship with the optimal policy in the case where the distribution of offers is exponential with the same mean as F.

In Sections 2 and 3 we consider a Bayesian version of this model by supposing that F is known to be one of the distributions $F_1, F_2, \ldots, F_n$ with given initial prior probabilities. In Section 2 we do not allow, and in Section 3 we do allow, the recall of old offers. In both cases we provide bounds on the optimal policy in terms of the optimal policies in the case where it is known which of the F_i is equal to F. This Bayesian format has previously been considered where it was assumed that F was a normal random variable with known variance and imposed a normal prior distribution on the mean of F. As our model imposes no parametric condition of F in the prior distribution, the type of results we obtain are somewhat different.